Enhancing Workspace Awareness on Collaborative Transparent Displays

Jiannan Li, Saul Greenberg and Ehud Sharlin
Department of Computer Science, University of Calgary
2500 University Drive NW, Calgary, AB, Canada
chrisleeseu@hotmail.com, saul.greenberg@ucalgary.ca, ehud@cpsc.ucalgary.ca

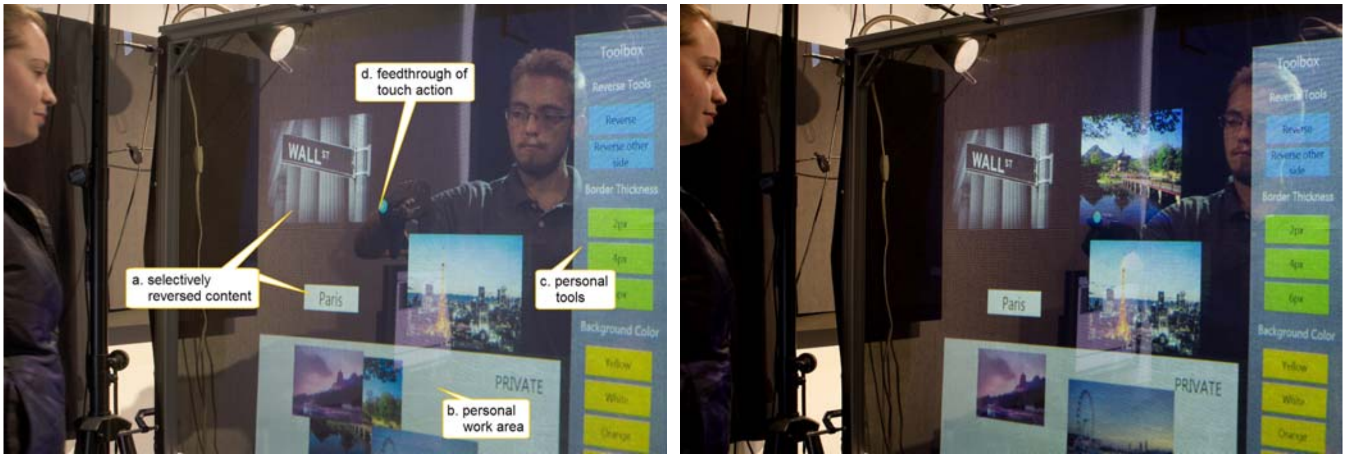

Figure 1. Our collaborative 2-sided transparent display. Note how transparency is compromised by graphics density.

ABSTRACT

Transparent displays can be used to support collaboration, where collaborators work on either side while simultaneously seeing what the other person is doing. This naturally supports *workspace awareness*: the up-to-the-moment understanding of another person's interaction with a shared workspace. The problem is that the transparency of such displays can change dynamically during a collaborative session, where it can degrade as a function of the density and brightness of the displayed graphics and changes in lighting. This compromises workspace awareness. Our solution is to track and graphically enhance a person's touch and gestural actions to make the *feedthrough* of those actions more visible on the other side. We had subjects perform three tasks over degrading transparency conditions, where augmentation techniques that enhance actions were either present or absent. Our analysis confirms that people's awareness is reduced as display transparency is compromised, and verifies that augmentation techniques can mitigate this awareness loss.

Author Keywords

Two-sided interactive transparent displays; workspace awareness, touch and gesture enhancement, CSCW.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI).

Cite as:

Li, J., Greenberg, S. and Sharlin, E. (2014) Enhancing Workspace Awareness on Collaborative Transparent Displays. Research report 2014-1065-16, Department of Computer Science, University of Calgary, Calgary, Alberta, Canada, October

INTRODUCTION

Transparent displays are 'see-through' screens. The basic idea is that a person can simultaneously view the graphics on the screen, while still seeing the real-world on its other side. One use of transparent displays is to support face-to-face collaboration [10, 9, 11], as such displays ostensibly provide two benefits 'for free'. As Figure 1 left illustrates, when a person is interacting on one side of a transparent screen, the person on its other side can see that person's gaze, hand and body movements *through* the display, as well as the changing graphics *on* the display. Seeing people's bodily actions relative to the artifacts in the workspace is critical for efficient collaborative interaction, as it helps communicate and coordinate mutual understanding. Technically, this is known as *workspace awareness*, defined as the "up-to-the-moment understanding of another person's interaction with a shared workspace" [5] (to be discussed shortly in detail).

For example, our own two-sided transparent display (Figure 1, left) allows people on either side to simultaneously interact with its projected graphics via touch and gestures [10]. It was also one of the first to afford different graphical contents on either side (see how [11] used fog), which we believe is important for several reasons. As annotated in Figure 1-a, text and image regions can be selectively reversed so that people on either side can view that content in its correct orientation. Next, (1-b,c) it affords personal work areas and tools that – while physically located on the same screen region – allows different content and interactions per side. Third, it means that visual *feedback* for the person performing an action on one side can differ from the visual *feedthrough* of that action as seen by the viewer on the other side, which can be important for particular collaborative situations [6]. For example, 1-c shows feedthrough of a person's touch action via a gold circle [10], which makes that touch visually prominent.

Yet our experiences with our own and other transparent displays revealed a critical problem: transparent displays are not always transparent [10]. All trade off the clarity of the graphics displayed *on* the screen *vs.* the clarity of what people can see *through* the screen. While transparency is partially inherent in the display technology, transparency also changes dynamically as a function of display content and external lighting (discussed shortly). This compromises what people can see and can severely affect workspace awareness. For example, compare Figure 1 right *vs.* left. A new photo placed in the center of the screen now makes the portion of the other person's body behind that image more difficult to see.

One solution to degraded transparency, and the subject of this paper, is to enhance feedthrough by tracking and visually augmenting human actions. Specifically, we explored two augmentation methods that can be easily applied to transparent displays. *Touch augmentation* highlights the current location of a fingertip, where a glow of increasing intensity and size is drawn on the other side of the display as the fingertip approaches the display, and that glow changes color when a touch is detected (Figure 1). *Trace augmentation* (inset & Figure 2) is somewhat similar, except a fading trace is drawn that follows the motion of the fingertip in space [10, 3, 4].

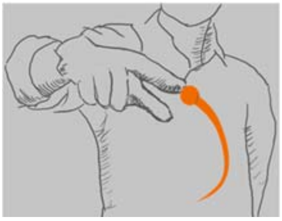

The question is, are touch and trace augmentation effective in supporting workspace awareness under degrading transparent display conditions? To answer this question, we conducted a study that investigated how people performed various collaborative tasks through a display. Participants performed three tasks under four different transparency levels (from highly transparent to barely transparent) where touch or trace augmentation methods were either present or absent. Our results show augmentation is highly effective when transparency is compromised, and incurs no penalty when transparency is uncompromised. The companion video figure illustrates this study. Before describing our study and our results, we begin with relevant background.

BACKGROUND

Workspace awareness

When people work together over a shared visual workspace (a large sheet of paper, a whiteboard, a touch display), they see both the contents and immediate changes that occur on that surface, as well as the fine-grained actions of people relative to that surface. This up-to-the-moment understanding of another person's interaction within a shared setting is the *workspace awareness* that feeds effective collaboration [5]. Workspace awareness provides knowledge about the 'who, what, where, when and why' questions whose answers inform people about the state of the changing environment: Who is working on the shared workspace? What is that person doing? What are they referring to? What objects are being manipulated? Where is that person specifically working? How are they performing their actions? In turn, this knowledge of workspace artifacts and a person's actions comprise key elements of *distributed cognition*: how cognition and knowledge is distributed across individuals, objects, artefacts and tools in the environment during the performance of group work [7].

People achieve workspace awareness by seeing how the artifacts present within the workspace change as they are manipulated by others (called *feedthrough*), by hearing others verbally shadow their own actions, by watching the gestures that occur over the workspace (called *intentional communication),* and by monitoring information produced as a byproduct of people's bodies as they go about their activities (called *consequential communication)* [5, 3, 4, 6].

Feedthrough and consequential communication occur naturally in the everyday world [5]. When artifacts and actors are visible at full fidelity, both give off information as a byproduct of action that can be consumed by the watcher. Thus consequential communication includes *gaze awareness*¸ where one person is aware of where the other is looking, and *visual evidence,* which confirms that an action requested by another person is understood by seeing that action performed. Similarly, intentional communication involving the workspace is easy to achieve in our everyday world. It includes a broad class of gestures, such as *deixis* where a pointing action qualifies a verbal reference (e.g., 'this one here') and *demonstrations* where a person demonstrates actions over workspace objects.

Workspace awareness plays a major role in various aspects of collaboration over a shared workspace [5].

- *Managing coupling.* People often shift back and forth between loosely-coupled mostly individual work, to tightly-coupled collaborative work. Awareness both enables and helps people perform these transitions.
- *Simplification of communication.* Because people can see the non-verbal actions of others, dialogue length and complexity is reduced.
- *Fine-grained coordination of action* is facilitated because one can see exactly what others are doing. This includes who accesses particular objects, handoffs, division of labor, how assistance is provided, and the interplay between peoples' actions as they pursue a simultaneous task.
- *Anticipation* occurs when people take action based on their expectations or predictions of what others will do. Consequential communication and outlouds play a large role in informing such predictions. Anticipation helps

people either coordinate their actions, or repair undesired actions of others before they occur.
- *Assistance.* Awareness helps people determine when they can help others and what action is required. This includes assistance based on a momentary observation (e.g., to help someone if one observes the other having problems performing an action), as well as assistance based on a longer-term awareness of what the other person is trying to accomplish.

Workspace awareness support in remote collaboration

In the late 1990s, various researchers in computer-supported cooperative work (CSCW) focused their attention on how distance-separated people could work together over a shared digital workspace. They quickly realized that early systems that showed only the shared graphics were insufficient. Because the partner could not see the other person's body, both intentional gestural communication and consequential communication was unavailable.

To overcome this, several researchers recreated face to face interaction via a 'see-though' display, typically done by blending a video of the remote person (or that person's silhouette) into the shared workspace [13, 14, 8]. This created the illusion that the geographically distant collaborators were on different sides of a transparent display, where one participant could see the artifacts as well as the remote participant on their screen.

Another strategy tracks a person's movements, and uses that information to graphically communicate that movement in the workspace as feedthrough. For mouse-based systems, multiple *telepointers* make each person's cursor visible to all. Telepointers become a surrogate for gestural actions, and suggest where that person is looking (gaze awareness) [2]. Telepointers can be augmented by *visual traces*, which visualize the last few moments' of a remote pointer's motion as a fading trail [3, 4]. For touch-based systems, the arms of multiple people working on either side can be digitally captured, where they are redrawn on the remote display in forms ranging from the realistic to abstract [12, 1]. What ties these and other methods together is the key idea that shared workspace technologies must recreate, as feedthrough, the otherwise lost cues of how the other person is interacting with the workspace.

Our work is similarly concerned with workspace awareness enhancements that facilitate how a person 'sees through' the display to view the person and their actions on other side. It differs in that we focus on collocated collaboration, where the display's transparency may be intermittently compromised during a collaborative session.

Factors Affecting Display Transparency.

Various factors interact to affect display transparency.

- *Graphics display technology.* Different technologies vary greatly in how they draw graphics (e.g., pixels) on a transparent display, e.g., dual-sided projector systems [10, 11], OLED and LCD screens, and even LEDs moving at high speed [9]. These interact with other factors to affect how people see through the screen.
- *Screen materials* can afford quite different levels of translucency, where what one sees through the display is attenuated by the material used [e.g., 9, 10, 11]. For example, manufactured screens sandwich emissive and conductive layers between glass plates in OLED displays, which affects its transparency. Our own work uses fabric with large holes in it as the screen material: the tradeoff is that larger holes increase transparency, while smaller holes increase the fidelity of the displaying graphics (Figure 1) [10].
- *Graphics density.* A screen full of high-density, busy, and highly visible graphics compromises what others can see through those graphics. That is, it is much harder to see through cluttered (*vs.* sparse) graphics (e.g., Figure 1 right *vs.* left).
- *Brightness*. It is harder to see through screens with significant bright, white (*vs.* dark) content, particularly if graphics density is high. Somewhat similarly, bright projectors can reflect back considerable light, affecting what people see through it.
- *Environmental lighting.* Glare on the screen as well as lighting on the other side of the screen can greatly affect what is visible through the screen. Similarly, differences in lighting on either side of the screen can produce imbalances in what people see. This is akin to a lit room with an exterior window at night time: those outside can see in, while those inside see only their own reflections.
- *Personal lighting.* If people on the other side of the display are brightly illuminated, they will be much more visible through the display than if they are poorly lit.
- *Clothing and skin color* and their reflective properties can affect a person's visibility through the display. Figure 1, for example, show the person on the other side wearing a black shirt and black glove, which negatively affects the visibility of his hand, arm and torso. In contrast, the bare hand seen in Figure 2 is much more visible. A white reflective glove would be even better.

Because of these factors, transparency (and thus the visibility of the other person) can alter dramatically throughout a collaborative interactive session. Screen materials and graphics display technology are static factors, but all others are dynamic. Graphics density and brightness can change moment by moment as a function of screen content. Lighting changes by shadows, by interior lighting turned on and off, and by the exterior light coming into the room (e.g., day *vs.* nighttime lighting). Clothing, of course, will vary by the person.

STUDY METHODOLOGY

Our study concerns itself with the interplay between transparency and workspace awareness. For terminology convenience, the *viewer* is the person (the participant) who observes the *actions* of the *actor* (the experimenter) on the

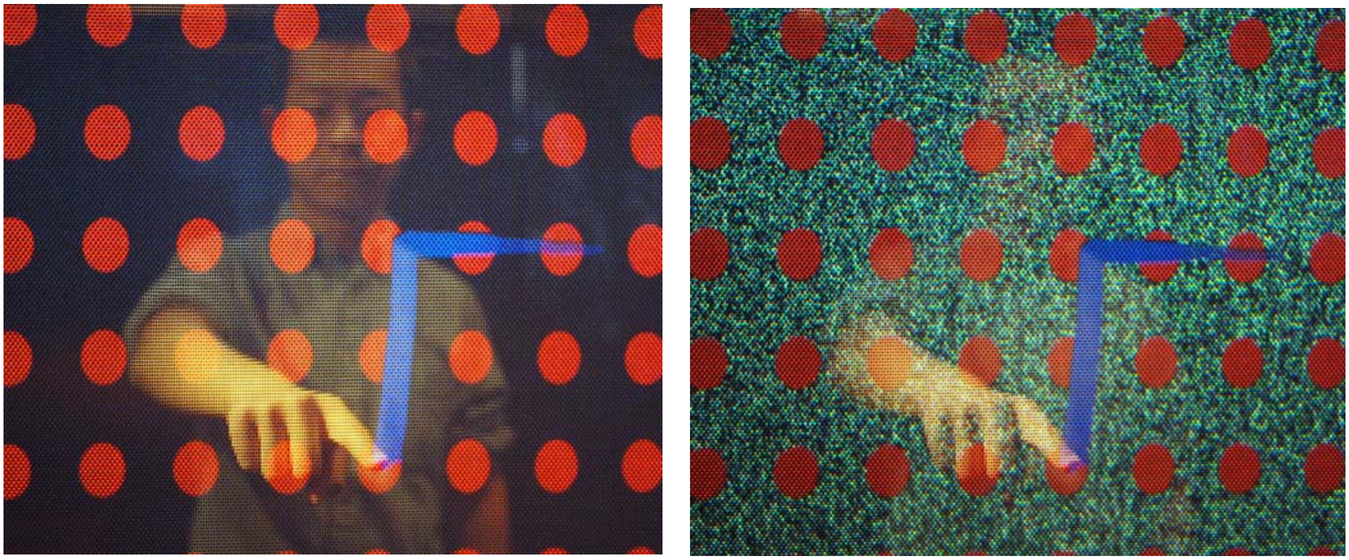

Level 1 transparency / front lit actor (actor clearly visible)

Level 2 transparency (body somewhat visible, hand visible

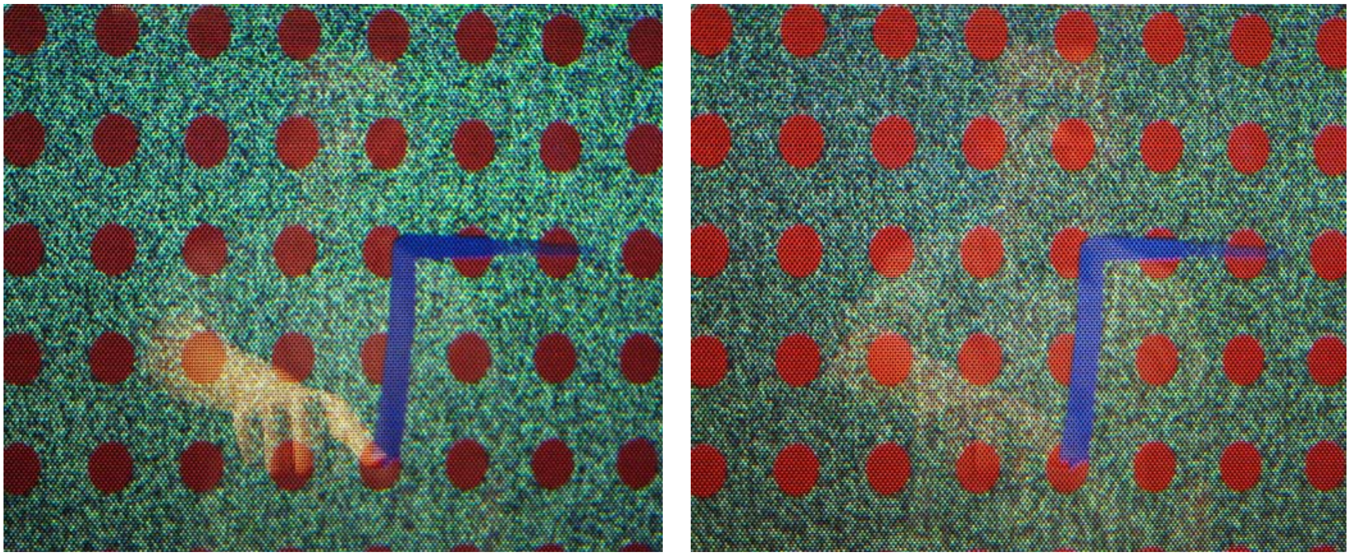

Level 3 transparency (body barely visible, hand somewhat visible)

Level 4 transparency (body / hand barely visible)

Figure 2. The 4 transparency conditions with trace augmentation on. All show the actor tracing a route (route task).

other side of the display. Our first hypothesis is that viewer's workspace awareness degrades as transparency is compromised. Our second hypothesis is that this degradation can be mitigated by enhancing the actor's actions via touch and trace augmentation methods.

Independent Variables

Transparency. We vary transparency as an independent variable. We use four transparency levels, each comprising a particular mix of graphical density patterns (projected onto the viewer's side of the display) and actor lighting. To explain, Figure 2 illustrates the 4 transparency conditions[1]. As will be explained shortly, all sub-figures show the actor in the same pose indicating a route through several circles, with trace enhancement turned on. The actor in all but the bottom right is front-lit. At the top left of Figure 2 is *level 1*, the most transparent condition, where the actor's hand, arm, body and eye gaze are clearly visible through the display. The top right is *level 2*, where we increase the graphical density by projecting a pseudo-random pattern comprising a ratio of 25% white to black pixels[2]. The actor's arm and hand are still clearly visible, but details of his body and eye gaze are harder to make out. The bottom left is *level 3*: the ratio

---

[1] To make images print-legible, we altered the lighting somewhat from the actual experimental conditions, and portray the actor in Figure 2 without gloves. However, the images are reasonable approximations of what study participants saw.

[2] We use an artificial pattern instead of photographs and text (in contrast to Figure 1), as we wanted to control transparency across the entire screen by creating a uniform wash.

is 67% and the actor's details become even more difficult to see (although the hand remains reasonably visible). The bottom right is *level 4*: the ratio remains at 67% but the actor is no longer front-lit. Here, the actor – while still discernable - is barely visible.

Augmentation: Enhancing Touch and Gestures. We developed two feedthrough augmentation techniques that try to enhance the viewer's visibility of the actor's touch and gestural actions [10]. As previously explained, the *augmented touch technique* draws a circular glow on the screen location corresponding to the actor's finger. The glow becomes larger and visually more intense as the actor's finger approaches the display, where the glow changes color when the display is actually touched (Figure 1 left). The *augmented trace technique* draws a fading line on the display, where the line follows the path of the actor's finger (Figure 2 inset). We treat augmentation as an independent variable, where it is either present or absent. The particular augmentation technique used (touch vs trace) depends upon the particular task associated with each study.

Tasks and Dependent Variables

We developed three tasks that exemplify common activities that people may perform on a two-sided display, where our tasks are variations of those describe in [4]. As mentioned, the experimenter is the actor, while the participant is the viewer. The viewer's performance over these tasks in our 8 conditions are our dependent variables, where they serve as a measure of their ability to maintain workspace awareness.

The shape task / error rate. *Shape gestures* refer to finger movements that trace geometric shapes that convey symbolic meanings, e.g., a character, a rightwards gesture indicating direction. Shape gestures can appear anywhere, and are not necessarily associated with workspace artifacts.

The *shape task* involves shape gesture actions. The actor uses his finger to 'write', as a shape gesture, a horizontally-reversed English letter over a randomly selected quadrant just above the display surface (reversal correctly orients the letter to the viewer). The viewer's task was to say out loud the letter s/he saw. We note that this task also required the viewer to disambiguate those parts of the gesture that were not part of the letter (e.g., when the person's finger approached and left the display surface). For augmentation conditions, we use the trace augmentation technique.

Error rate is the dependent variable: the number of incorrectly recognized or missed shapes over the total number of shapes presented per condition.

Route task / accuracy rate. *Route gestures* are paths going through some objects in the workspace. Routes can suggest actual paths in the space, transitions between object states, or groupings of objects. Unlike shape gestures, they are made relative to the workspace and its artifacts.

The *route task* involves route gesture actions. A 16x10 grid of circles are aligned to appear on the same locations on both the actor's and viewer's sides of the screen. The actor then gestures a path through a particular sequence of circles (illustrated in Figure 2). While routes differed between trials, all paths went through five circles with one turn in the middle. The viewer's task was to reproduce that path by touching the circles the path went through. We use the trace augmentation for the augmentation conditions.

Accuracy rate is the dependent variable: the number of correct responses over the total number of responses per conditions. Correct responses are those that state all circles the gesture went through.

The point task / response time, response error, miss rate. The previous tasks are examples of tightly-coupled collaboration: both actor and viewer focus their attention on the gesture as it is being performed. We wanted to see what would happen in *mixed-focus collaboration*, where participants pursue individual work while still monitoring group activities [6, 5]. As previously mentioned, workspace awareness is particularly important for mediating the shift from loosely to tightly coupled group work, for it helps create opportunities to coordinate mutual actions.

The *point task* measures, in part, a viewer's ability to stay aware of the actor's touch actions during mixed-focus collaboration. The viewer, while performing individual work, had to simultaneously monitor the actor and indicate when s/he saw the actor touch the work surface. We use touch, as most contemporary interaction methods require the actor to touch the display to manipulate the workspace artefacts. The actor taps a randomly-positioned circle that appears only on his side of the display. That circle disappears, a new circle positioned elsewhere appears somewhat afterwards, and the process repeats. To emulate mixed-focus collaboration, the viewer had two tasks. For the *individual task,* the viewer was asked to tap solid squares as they appeared on the viewer's side of the display. In the *follower task*, the viewer was asked to tap those spots that s/he had noticed were touched by the actor. The viewer was told that the follower task took precedence over the individual task, where s/he had to react as quickly and as accurately as possible to indicate where the actor had touched. On average the ratio of individual to follower task episodes were ~3:1, but were interleaved irregularly to make their timing unpredictable to the viewer. We use the touch augmentation for the augmentation conditions.

Three metrics measured awareness as a dependent variable. *Response time* is the elapsed time between the touch from the actor and the following responding touch from the viewer. *Response error* is the distance between the location touched by the actor and the location touched by the viewer. *Miss rate* is the rate where participants failed to react to a touch by the actor, e.g., because the viewer didn't notice the touch or failed to see where the touch occurred.

Study Design

We ran three studies. Each study is similar in form, except that participants performed a different task (shape, route and point), each with their own dependent variables. All are based upon a within-subject (repeated measures) ANOVA factorial design: *transparency* (4 levels) x *augmentation* (2 levels), or 8 different conditions per task. All used the same participants as viewers, where each participant did all three tasks over all 8 conditions (with many repeated trials per condition) in a single 90 minute session. For each condition, subjects underwent many repeated trials. Transparency levels are as described above. Augmentation type varies per task, and is either present (augmentation on) or absent (augmentation off).

Hypotheses

Our null hypothesis is suggested by our study design.
There is no difference in participant's ability to
(a) recognize the shape as measured by the error rate,
(b) trace a route as measured by the accuracy, and
(c) observe touches as measured by the response time, the response error, and the miss rate,
across the four transparency levels and the presence or absence of augmentation.

Materials

The study was conducted on our two-sided transparent display prototype, with technical details described in [10]. In essence, it is a 57x36 cm two-sided transparent display, where projectors on each side project its visuals. An OptiTrack Flex 13 motion capture system tracked a marker placed on the index finger of gloves worn by participants. Dedicated software modules displayed screen contents for each task, and collected data about user actions.

Participants

Twenty-four participants (10 female and 14 male) between the ages of 19 and 41 were recruited from a local university for this study. All were experienced in some form of touch screen interactions (e.g., phones, surfaces). All were right-handed. Each participant received a $15 payment.

Procedure

After being briefed about the study purpose, the participant completed a demographics questionnaire. Participant then performed the shape, route and point task in that order. For each task, the participants were instructed on what they had to do, and then did 9 blocks: a practice block and then eight counter-balanced blocks corresponding to the eight previously described conditions. After completing each task, the experimenter led the participant through a semi-structured interview, where the participant was asked to comment about his or her experiences with the various conditions, as well as the strategies used to perform tasks.

RESULTS

Statistical Analysis Method

We ran a two-way repeated measures ANOVA for each of the measures obtained from the three tasks, with sphericity assumed. For sphericity-violated cases, we used Greenhouse-Geisser corrections. For post-hoc tests, we used the test of simple main effects with Bonferroni corrections. The level of significance was set at $p<0.05$.

The Shape Task

In the shape task, the actor wrote, as a gesture, a horizontally reversed capital letter; the viewer's task was to say what letter he or she saw. The error rate of the shape task was then calculated as the ratio of misrecognized letters in each condition for each participant.

Results. Our analysis reported a significant main effect for transparency (F3, 69 = 12.458, $p < 0.05$), augmentation (F1, 23 = 42.037, $p < 0.05$), and the interaction between them (F3, 69 = 14.73, $p < 0.05$).

Figure 3 graphically illustrates the means of the error rate and our post-hoc test results. The green and blue lines represent the augmentation on *vs.* off conditions respectively, while the four points on those line are the values measured at each of the four transparency levels, with level 1 on the left and level four on the right. The vertical red lines indicate where the post-hoc test reported a significant difference between the augmentation off *vs.* on values at a particular transparency level. For example, we see that the red lines indicate a significant difference in the error rate between the augmentation on/off conditions at levels 2, 3 and 4. The numbers in the colored box next to particular points indicate which transparency levels differed significantly on a given augmentation condition. For example, with augmentation off, we see from the numbers in the blue box that: level 1 differs significantly from levels 3 and 4; and levels 2 and 3 differ from level 4. However, with augmentation off, there are no significant differences in the error rate at any transparency level.

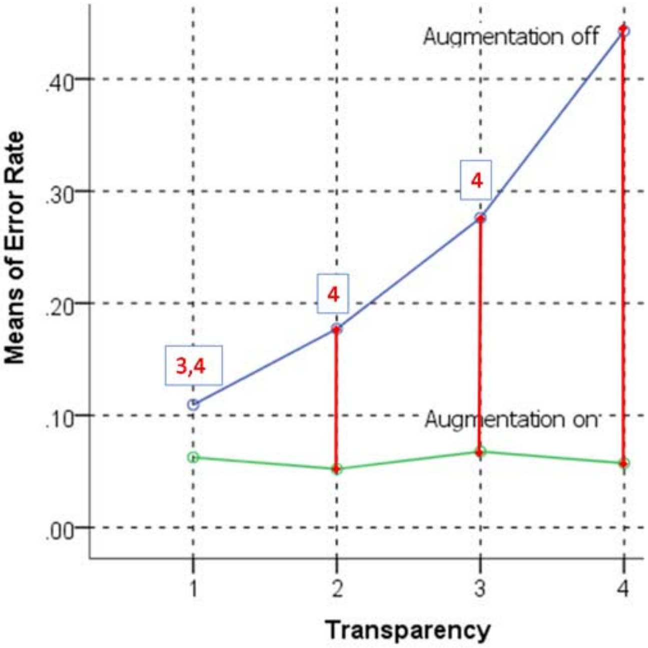

Figure 3. Shape task results. *Error rate* plotted by condition

Discussion. The null hypothesis for the shape task is rejected. First, without augmentation, there is a notable increase in the error rate as display transparency decreases, where most pairwise differences between these means are statistically significant (Figure 3, blue line). Differences are practically significant as well, where the error rate of ~10% in the most transparent condition increases to ~44% in the least transparent condition (see the blue line data points in Figure 3).

Second, with augmentation, the error rate is constant regardless of the transparency level, with no significant difference seen across any of the transparency levels when augmentation is used (Figure 3, green line). Notably, the error rate is low at ~6%. This sharply contrasts with augmentation off conditions, where the error rate increases as transparency decreases.

Third, the presence or absence of augmentation does not affect error rate in highly transparent conditions, i.e., using augmentation when it is not needed does not incur a negative effect (compare the first points in Figure 3's green *vs.* blue lines, where differences are not significant).

In summary, the results indicated that people have much more difficulty correctly recognizing shape gestures as transparency is compromised (without augmentation). They also indicate that the trace augmentation mitigates this problem, where people are able to maintain a largely stable and fairly low error rate (M = 6.0%, SD = 0.013) that is equivalent to highly transparent conditions. That is, the trace augmentation supports people's ability to perceive the other's gestural shapes as transparency deteriorates.

The Route Task

In the route task, the actor gestured a path through a particular sequence of circles shown on the display. The viewer's task was to reproduce the path by touching particular circles that the path went through. The accuracy of the route task was then calculated as the ratio of correctly reproduced paths to the total paths in each condition.

Results. Our analysis discovered a significant main effect for transparency (F3, 69 = 7.240, $p < 0.05$), augmentation (F1, 23 = 42.037, $p < 0.05$), and the interaction between them (F3, 69 = 4.515, $p < 0.05$). Figure 4 graphically illustrates the means of the accuracy rate and our post-hoc test results, where their portrayal is similar to Figure 3.

Discussion. The null hypothesis for the route task is rejected. First, without augmentation the accuracy decreases noticeably as display transparency deteriorates (Figure 4, blue line), where we see statistically significant differences between the accuracy at transparency level 1 and all other levels. The differences are also practically significant: the ~91% accuracy in the most transparent condition degrades to ~62% in the least transparent condition.

Second, accuracy across transparency levels in augmentation-on conditions is constant at a high level (~85-90%): the slight downward trend is not significant (Figure 2, green line). For transparency level 4, accuracy is significantly higher with augmentation than without.

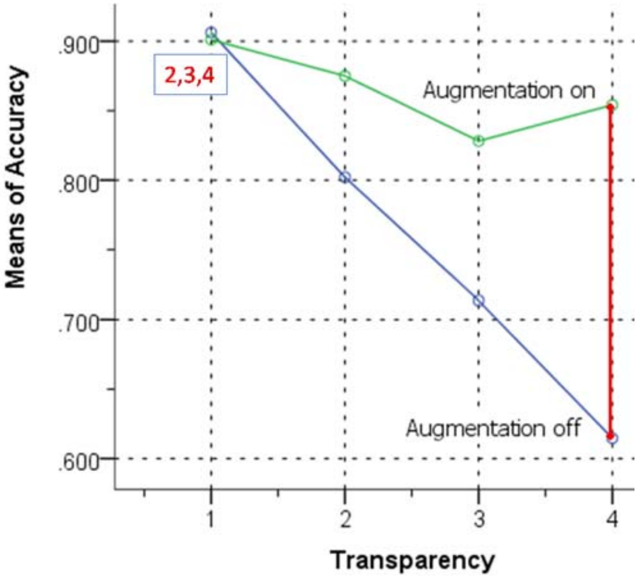

Figure 4. Route task results. *Accuracy rate* plotted by condition

Third, the presence or absence of augmentation does not affect accuracy in highly transparent conditions, i.e., it does not incur a negative effect (compare 1st points in Figure 4's green *vs.* blue lines, where differences are not significant).

To sum up, the results indicate that people have much more difficulty accurately perceiving the route gesture when display transparency is compromised (without augmentation). The results also indicate that trace augmentation alleviates these difficulties at low levels of transparency. That is, the trace augmentation supports people's ability to perceive the other's path drawing gestures relative to objects as transparency deteriorates.

The Point Task

In the point task, the viewer was asked to: (a) carry out a separate independent task, and (b) simultaneously monitor and respond to the actors' touch actions on the display by touching the location where the actor had just touched. Response time is the average elapsed time between the actor's touch and the responding viewer's touch. Response error is the distance between the location touched by the actor and the corresponding location touched by the viewer. Miss rate is the rate where viewers failed to react to the actor's touch.

Results: Response Time. Our analysis revealed a significant main effect for response time for transparency (F3, 69 = 20.731, $p < 0.05$), augmentation (F1, 23 = 4.517, $p < .05$), and the interaction between them (F3, 69 = 4.620, $p < 0.05$). Figure 5a graphically illustrates the means of the response time and our post-hoc test results.

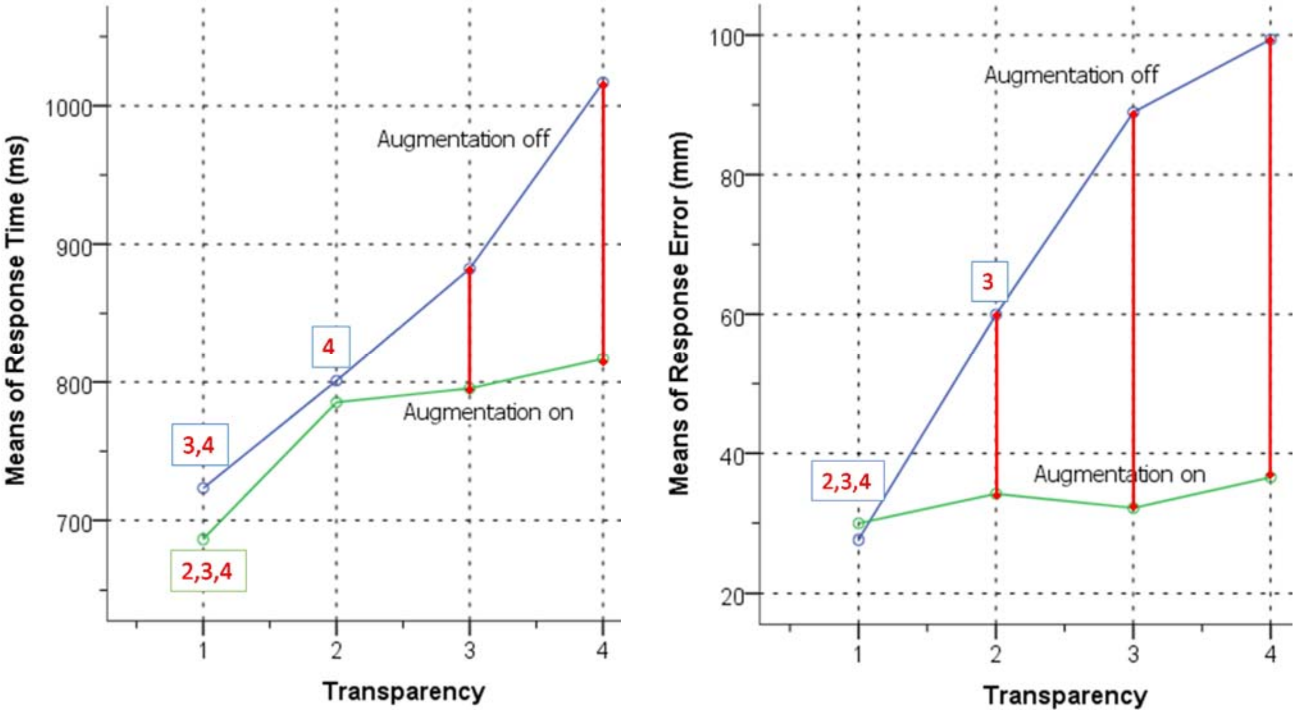

a) *response time* by condition

b) *response error* by condition

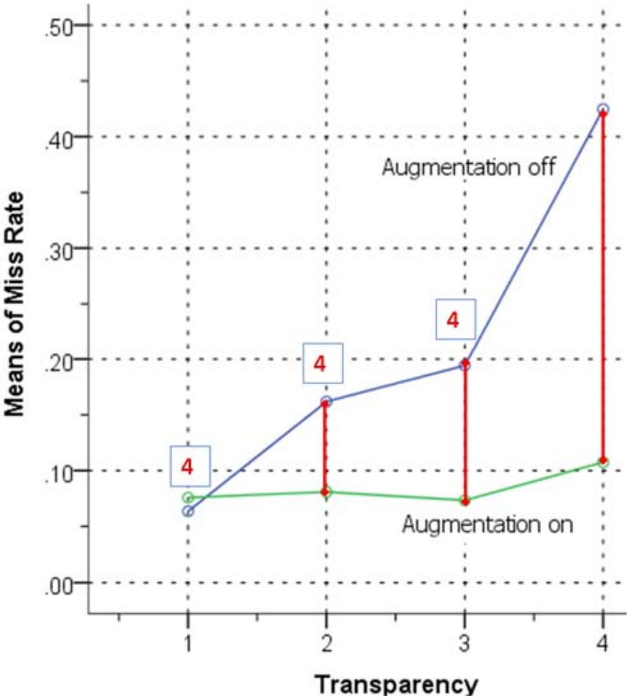

c) *miss rate* by condition

Figure 5. Point task results.

Discussion: Response Time. The null hypothesis is rejected. First, without augmentation, response time tends to increase as display transparency decreases (significant differences are visible between these means in Figure 5a, blue line). The differences are also practically significant, with response times of ~700ms increasing to ~1000ms between the most to least transparent conditions.

Second, with augmentation the response time exhibits a statistically significant but somewhat modest increase from transparency level 1 (~700ms) to level 2 (~800ms), with no further increase afterwards (Figure 5a, green line).

Third, for levels 1 and 2 transparency, adding augmentation neither increases nor reduces the response time with respect to similar conditions without augmentation i.e., it does not incur a negative effect. Yet augmentation is beneficial in low transparency conditions (compare Figure 5a data points between the green and blue lines).

In summary, the results indicate that people pursuing their own individual tasks while simultaneously monitoring another person's touches are somewhat slower to respond when transparency is compromised (without augmentation). The results also indicate that the touch augmentation method mitigates this somewhat: their response time increases only slightly in low transparency conditions.

Results: Response Error. Our analysis revealed a significant main effect on response error for transparency (F3, 69 = 11.676, $p < 0.05$), augmentation (F1, 23 = 48.508, $p < 0.05$), and the interaction between them (F3, 69 = 13.270, $p < 0.05$). Figure 5b graphically illustrates the means of the response error and our post-hoc test results.

Discussion: Response Error. The null hypothesis is rejected. First, without augmentation the response error increases as display transparency deteriorates (significant differences are visible between these means in Figure 5b, blue line). The differences are also practically significant, where the response error of ~28mm in the most transparent condition increases threefold to ~99mm in the least transparent condition.

Second, with augmentation the response error is constant regardless of the transparency levels, with no significant differences between them (Figure 5b, green line). Furthermore, the response error stays low (at ~33mm) when augmentation is present; this contrasts dramatically to the statistically significant increase in response error without augmentation when display transparency is compromised (compare green and blue lines in Figure 5b).

Third, the presence or absence of augmentation does not affect error rate in highly transparent conditions, i.e., it does not incur a negative effect. Yet it is beneficial in all other conditions when transparency is compromised (compare Figure 5b data points between the green and blue lines).

In summary, the results indicate that people are less precise when display transparency is compromised (without augmentation). The results also indicate that the touch augmentation method mitigates this considerably.

Results: Miss Rate. Our analysis found a significant main effect on the miss rate for transparency (F3, 69 = 23.249, $p < 0.05$), augmentation (F1, 23 = 21.300, $p < 0.05$), and the interaction between them (F3, 69 = 15.434, $p < 0.05$). Figure 5c graphically illustrates the means of the response time and our post-hoc test results.

Discussion: Miss Rate. The null hypothesis is rejected. First, without augmentation the miss rate increases sharply as transparency is reduced where a significant difference is seen between the first 3 levels *vs.* the 4th level (Figure 5c, blue line). This difference is practically significant, where the miss rate jumps from ~6% in the most transparent condition to ~43% in the least transparent condition.

Second, with augmentation the miss rate remained invariably low at ~8% (Figure 5c, green line).

Third, the presence or absence of augmentation does not affect error rate in highly transparent conditions, i.e., it does not incur a negative effect. Yet it is beneficial in all other conditions when transparency is compromised (compare Figure 5c data points between the green and blue lines).

In summary, the results indicate that people, when pursuing their own individual tasks while simultaneously monitoring another person's touches, are much more likely to miss the other person's touch actions when transparency is compromised (without augmentation). The results also indicate that the touch augmentation method mitigates this: the miss rate remains low under all transparency conditions.

Overall discussion of results

The above results, when considered collectively, consistently show that decreasing display transparency reduces a viewer's awareness of the actor's actions on the other side of a transparent display. Across all three tasks and as reflected by all five measures, participants' performance with no augmentation generally deteriorated as transparency was compromised. Differences were both statistically and practically significant.

The same results also show that augmentation techniques mitigate awareness loss when display transparency is compromised. Again, this was true across all tasks and all measures, where differences were both statistically and practically significant.

We also saw that the augmentation techniques did not have a negative effect in situations where they were not strictly necessary, i.e., high transparency conditions when the actor's actions are clearly visible. Across all tasks and for 4 of the 5 measures, the presence or absence of augmentation had little effect on participants' performance at the highly transparent level. On the other hand, we also saw that augmentation almost always had a beneficial effect when transparency was degraded when compared to the no-augmentation condition.

However, the results also reveal subtleties. While all measures in all tasks show that augmentation helps overcome the degradation in people's performance as transparency declines, it is not always continuous. For example, consider the response time measure in the point task, as illustrated in Figure 5a, where there is a difference between the response time in the augmentation on condition between levels 1 and 2. Thus we see an (isolated) case where workspace awareness has degraded, but augmentation does not appear to help. Our post-study interviews of participants suggest why this is so. Most reported that their strategy was to watch for movements of other body parts of the actor *before* the finger was close to the screen (e.g., raising the arm and moving the hand towards the screen). This consequential communication signaled that a touch was soon to occur. Participants said they found it increasingly difficult to see those body movements as transparency decreased, and consequently they reacted more slowly. For example, at transparency level 2 (Figure 2, upper right), people found it more difficult to see initial arm movements, but they could still see the hand as it approached the display. While touch augmentation provided information about where the fingertip was and its distance to the screen, it did not signal the earlier actions of other body parts and thus had no net benefit. When transparency was compromised even further at levels 3 and 4, participants had more difficulty seeing the un-augmented approaching finger (Figure 5a. blue line). In those cases, augmentation helped signal the approach at closer ranges, thus enabling people to react faster as compared to no augmentation (Figure 5a, green line).

Overall, we conclude that augmentation can supply the information necessary for people to maintain workspace awareness as transparency degrades. In those cases where augmentation may not provide any benefit (such as highly transparent situations where the actor is clearly visible), augmentation can still stay on as it has no negative effects. Keeping augmentation on at all times is useful, as our results also show that the degradation of workspace awareness varies (more or less) as a function of transparency degradation: there is no clear threshold that defines when augmentation should be turned on.

IMPLICATIONS

Providing necessary workspace awareness is crucial for the utility and usability of collaborative transparent displays. Therefore, their hardware and software interface design should guarantee reasonable support for the cues that comprise workspace awareness. We offer two implications for addressing this awareness requirement.

Implication 1: Controlling Transparency

Transparent displays are often portrayed as fully transparent in commercial advertisements, many research figures, and even futuristic visions of technology. We suspect that their graphics density and lighting are tuned to show such displays at their best. Yet transparent displays are not invariantly transparent. The consequence (as our results clearly show) is that degrading transparency can greatly affect how collaborators maintain mutual awareness.

One partial solution is to control display transparency as much as possible. Our experimental setup and study confirmed that high graphics density and dim lighting on the actor can reduce what one can see through the display. This can be partially remedied by design. For lighting, the system could incorporate illumination sources (perhaps integrated into the display frame) that brightly illuminates the collaborators. For graphics density, applications for transparent displays should distribute graphics sparsely on the screen, with enough clear space between its elements to permit one to see through those spaces. Colors, brightness and textures can be chosen to find a balance between seeing the displayed graphics and seeing through those graphics.

Another partial solution controls for external factors. This includes the ambient light that may reflect off the display, and even the color of surrounding walls and furniture. For example, we surrounded our own display with blackout curtains both to block out light and to provide a dark background [10]. Another controllable factor is the color of the collaborators' clothes (bright colors are more reflective than dark colors) and how that color contrasts with the surrounding background. For example, participants can wear white reflective gloves to better illuminate their hand movements to others.

Another partial solution relies on the display technology itself. For example, our display is based on a mesh fabric that only allows a certain amount of light to pass through it [10]. Other technologies, such as JANUS [9], may afford more light transmission. However, we should not expect technical miracles, as we believe that all technologies will be affected by the factors mentioned earlier in this paper.

In practice, we expect that the ability to control for the above factors is highly dependent on context. Designers may be able to devise (or recommend) specific transparency modulation mechanisms if they know where the display is used and what tasks people are carrying out on it. However, we expect most installations will limit what designers can control. Fortunately, we can still enhance workspace awareness by augmenting user actions, as discussed next.

Implication 2: Augmenting User Actions

Our study revealed that augmentation techniques can mitigate awareness loss when display transparency is compromised. In spite of the simplicity of our techniques (revealing the motion of a single finger), they proved effective. This clearly suggests that – at the very least – designers should visually augment a person's dominant finger movements. This is somewhat generalizable, as that finger often signals pointing gestures, is the focal point of input interaction for touch-based displays, and hints at where the actor is directing their gaze.

However, we can do even better. While seeing finger movement is helpful, body language is far richer. In daily face-to-face activities, we maintain workspace awareness by observing movements of multiple body parts (including gaze awareness) and interpret those sequences in relation to the workspace. We need to develop augmentation techniques that capture that richness, where we expect it will be helpful across a broader variety of tasks and situations. Examples include systems that: represent the entire hand, that change the representation as a function of distance; that show where a person is looking; that show the entire arm [12], or even that show the entire body [14].

Of course, there are challenges to this. Technical challenges include tracking. Graphical challenges include designing an easily understood representation that does not occlude, distract, or otherwise interfere with a person's view of the workspace: recall that workspace awareness involves a view of the participant, the workspace artifacts, and the participant's actions relative to those artifacts.

In summary, simple augmentation techniques will likely work well for mitigating awareness loss in many scenarios. However, new techniques and representation should be developed to better match the situation, display and task.

LIMITATIONS

Our controlled study was, to our knowledge, the first of its kind and, as typical with such studies, has limitations.

First, we used only four transparency levels. While these were chosen to capture a range from highly to barely transparent, it does not cover the full transparency spectrum nor expose other factors that could affect transparency.

Second, our manipulation of graphical density was artificial, where we used a random pixel pattern containing a well-defined ratio of bright *vs.* dark pixels as a wash. Real world graphics are different, where we could have tested how people maintain awareness through (say) a document editor, a photo-viewing application, and/or a running video.

Third, the three study tasks were artificial. They cover only a small set of tracing gestures and touch actions that people perform during cooperative work. Our augmentation methods matched what we thought would be critical actions. While we consider these tasks reasonable representatives of what people do during collaboration, they do not cover all interaction nuances. As well, the tasks did not test people doing real tasks, where people may exhibit more complex interaction and gestural patterns.

CONCLUSION

Our study investigated the effect of display transparency on people's awareness of others' actions, and the effectiveness of augmentation techniques that visually enhance those actions. Our analysis confirms that people's awareness is reduced when display transparency is compromised, and that augmentation techniques can mitigate awareness loss. Based on our findings, we suggested a few implications for collaborative transparent display designers.

ACKNOWLEDGMENTS

Funds provided by NSERC-AITF-SMART Industrial Chair in Interactive Technologies, NSERC's Discovery Grant and Surfnet Network. Sutapa Dey helped in our pilot studies.

REFERENCES

1. Genest, A., Gutwin, C., Tang, A., Kalyn, M. and Ivkovic, Z. (2013) KinectArms: a Toolkit for Capturing and Displaying Arm Embodiments in Distributed Tabletop Groupware. *Proc. ACM CSCW,* 157-166.
2. Greenberg, S., Gutwin, C. and Roseman, M. (1996) Semantic Telepointers for Groupware. *Proc. OZCHI*, IEEE Press, 54-61.
3. Gutwin, C. (2002) Traces: Visualizing the Immediate Past to Improve Group Interaction. *Proc. Graphics Interface,* 43-50.
4. Gutwin, C. and Penner, R. (2002) Improving Interpretation of Remote Gestures with Telepointer Traces. *Proc. ACM CSCW,* 49-57.
5. Gutwin, C. and Greenberg, S. (2002) A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *J. CSCW,* 11(3-4):411-446.
6. Gutwin, C. and Greenberg, S. (1998) Design for Individuals, Design for Groups: Tradeoffs Between Power and Workspace Awareness. *Proc. ACM CSCW.*
7. Hollan, J., Hutchins, E. and Kirsh. D. (2000) Distributed Cognition: Toward a New Foundation for Human-Computer Interaction Research. *ACM TOCHI*, 7(2).
8. Ishii, H. and Kobayashi, M. (1992) ClearBoard: a Seamless Medium for Shared Drawing & Conversation with Eye Contact. *Proc. ACM CHI*, 525-32.
9. Lee, H., Hong, J., Lee, G. and Lee, W (2014) Janus. *ACM SIGGRAPH 2014 Emerging Technologies,* #15.
10. Li, J., Greenberg, S., Sharlin, E. and Jorge, J. (2014) Interactive Two-Sided Transparent Displays: Designing for Collaboration. *Proc. ACM DIS*, 395-404.
11. Olwal, A., DiVerdi, S., Rakkolainen, I. and Hollerer, T. (2008) Consigalo: Multi-user Face-to-face Interaction on Immaterial Displays. *Proc. INTETAIN*, #8, ICST
12. Tang, A., Boyle, M. and Greenberg, S. (2004). Display and Presence Disparity in Mixed Presence Groupware. *Proc. Australasian User Interface Conf* (Vol. 28. Australian Computer Society, Inc., 73-82.
13. Tang, J. and Minneman, S. (1990) Videodraw: A Video Interface for Collaborative Drawing. *Proc ACM CHI.*
14. Tang, J. and Minneman, S. (1991) VideoWhiteboard: Video Shadows to Support Remote Collaboration. *Proc. ACM CHI*, 315-322